

THE INVENTION CLAIMED IS:

1. A hierarchical method for fault tolerance in a distributed computer system:
providing a plurality of data centers;
providing a plurality of objects in each of the plurality of data centers;
5 using a local sub-protocol for dissemination of messages within a data center in the plurality of data centers; and
activating the local sub-protocol from another data center of the plurality of data centers in a single round-trip message in the absence of faults.
2. The hierarchical method as claimed in claim 1 wherein:
10 using the local sub-protocol uses an atomic broadcast protocol.
3. The hierarchical method as claimed in claim 1 wherein:
using the local sub-protocol uses an atomic broadcast protocol and invokes the plurality of objects in the data center.
4. The hierarchical method as claimed in claim 1 wherein:
using the local sub-protocol uses an atomic broadcast protocol and invokes the plurality of objects in other of the plurality of data centers by sending propagation messages;
and including:
20 responding to the propagation message in the other of the plurality of data centers activates a local atomic broadcast protocol.
5. The hierarchical method as claimed in claim 1 wherein:
using the local sub-protocol uses an atomic broadcast protocol and invokes the plurality of objects in other of the plurality of data centers by sending a propagation message;
25 and including:
responding to the propagation message in the other of the plurality of data centers includes providing an acknowledgement to the data center in the plurality of data centers from one of the plurality of objects therein.
6. The hierarchical method as claimed in claim 1 wherein:
30 using the local sub-protocol uses an atomic broadcast protocol and invokes the plurality of objects in other of the plurality of data centers by sending a propagation message;

and including:

responding to the propagation message in the other of the plurality of data centers includes providing an acknowledgement to the data center in the plurality of data centers from one of the plurality of objects therein; and

5 waiting a time for the acknowledgement and sending a second propagation message to another of the plurality of objects in the other of the plurality of data centers if the acknowledgement is not received within the time.

7. The hierarchical method as claimed in claim 1 wherein:

10 using the local sub-protocol uses an atomic broadcast protocol and invokes the plurality of objects in other of the plurality of data centers by sending a first propagation message;

and including:

responding to the first propagation message in the other of the plurality of data centers includes providing an acknowledgement to the data center in the plurality of data centers from one of the plurality of objects therein;

waiting a time for the acknowledgement and activating a second propagation message to another of the plurality of objects in the other of the plurality of data centers if the acknowledgement is not received within the time; and

20 sending the first and second propagation messages includes sending first and second unique identifiers.

8. The hierarchical method as claimed in claim 1 wherein:

activating the local sub-protocol includes using a unique identifier.

9. The hierarchical method as claimed in claim 1 wherein:

25 providing the plurality of objects includes providing a primary object in one of the plurality of data centers that communicates with the other of the plurality of data centers.

10. The hierarchical method as claimed in claim 1 wherein:

using the local sub-protocol includes detecting failures within the plurality of objects within each of the plurality of data centers.

30 11. The hierarchical method as claimed in claim 1 wherein:

using the local sub-protocol includes determining when faults occur in the plurality of objects in a local data center.

12. The hierarchical method as claimed in claim 1 wherein:
using the local sub-protocol includes determining when faults occur in the plurality of
objects in a local data center;
and including:
5 determining an alternate object in the plurality of objects in the local data center to
become a backup primary object when a primary object has a fault.

13. The hierarchical method as claimed in claim 1 wherein:
using the local sub-protocol includes developing a suspicion of the occurrence of
10 faults in the plurality of objects in a local data center;
and including:
determining an alternate object in the plurality of objects in the local data center to
become a backup primary object when a primary object has a fault.

14. A hierarchical method for fault tolerance in a distributed computer system:
providing a plurality of data centers;
providing a plurality of objects in each of the plurality of data centers;
using a local sub-protocol including an atomic broadcast protocol for fault-tolerant
15 dissemination of messages within a data center in the plurality of data centers;
and
activating the local sub-protocol from another data center of the plurality of data
centers using a propagation message sent in a fault-tolerant manner in a single
round-trip message in the absence of faults.

15. The hierarchical method as claimed in claim 14 including:
20 responding to the propagation message in the other of the plurality of data centers
includes providing an acknowledgement to the data center in the plurality of
data centers from one of the plurality of objects therein.

16. The hierarchical method as claimed in claim 14 including:
25 responding to the propagation message in the other of the plurality of data centers
includes providing an acknowledgement to the data center in the plurality of
data centers from one of the plurality of objects therein; and
30 waiting a time for the acknowledgement and sending a second propagation message to
another of the plurality of objects in the other of the plurality of data centers if
the acknowledgement is not received within the time.

17. The hierarchical method as claimed in claim 14 including:
responding to the first propagation message in the other of the plurality of data centers
includes providing an acknowledgement to the data center in the plurality of
data centers from one of the plurality of objects therein;
waiting a time for the acknowledgement and sending a second propagation message to
another of the plurality of objects in the other of the plurality of data centers if
the acknowledgement is not received within the time; and
sending the first and second propagation messages includes sending first and second
unique identifiers.

18. The hierarchical method as claimed in claim 14 wherein:
activating the local sub-protocol includes using the propagation message with a
unique identifier.

19. The hierarchical method as claimed in claim 14 wherein:
using the local sub-protocol to determine when faults occur in the plurality of objects
in a local data center;

and including:
determining an alternate object in the plurality of objects in the local data center to become a backup primary object when a primary object has a fault.

20. The hierarchical method as claimed in claim 14 wherein:
using the local sub-protocol to developing a suspicion of the occurrence faults in the
plurality of objects in a local data center:

and including:
determining an alternate object in the plurality of objects in the local data center to become a backup primary object when a primary object has a fault